

Updated June 21, 2007

Q: ITAR does not seem to apply directly, but EAR is particularly relevant to the project. Reading through EAR Category 5-type 1 (just the telecommunications equipment part), I think we are ok for phase-1 since we are using commodity equipment (802.11) with quite limited bandwidth capabilities. However, Category 5-type 2 (information security) has me a little concerned. As I read this, it seems that any university that uses SSL is potentially in export violation.

We are basically planning on using open source (e.g. openssl) implementations of cryptographic functions. We have a very international team that has frequently published and used such algorithms. Will there be a problem here? If so, do you have any recommendations for us? For example, should we limit our encryption algorithms to 56bit algorithms? Or, can you explain to me how universities would, in general, be able to do such work?

A: It is the Company/University's responsibility to understand and adhere to the Government export control regulations with respect to the technologies they are using. If you are in doubt about the regulation of technologies you are using, or developing, you should consult your institution's legal council and/or the Departments of Commerce and State.

Reference: <http://www.bis.doc.gov/> ,
http://www.pmddtc.state.gov/itar_index.htm

Department of Commerce. Outreach and Educational Services Division: (202) 482-4811

Department of State, Directorate of Defense Trade Control Response Team Telephone Number:
(202) 663-1282

Response Team E-mail: DDTCResponseTeam@state.gov

We have used your query as a test case and engaged the resources above to find an answer. We found the system to be responsive and straightforward. You should have little difficulty finding the answers you seek and complying with the regulations.

Updated June 15, 2007

Q: The IAMANET PIP says “universities may submit by grants.gov” as opposed to “universities must submit”. May university-primed proposals be sent by fedex instead of the grants.gov process?

A: No, unfortunately not.

Q: When the PIP says to plan for 1 month of on-site evaluation by a red-team, does that mean that this month is part of phase 1, or is it a month in between phase 1 and a phase 2?

A: It is part of Phase 1

Q: Also, do we need to prepare a budget for red-team people to stay at our facility (e.g. hotels, meals, engineer time to conduct walkthrough, etc)?

A: The red team will be responsible for its own per diem expenses; the only expenses the performer teams need to consider is for on-site expenses.

Q: Please clarify the statement that the government will furnish baseline MANET protocols. Is that meant to imply, for example, that the government will supply working Linux-based code of OLSR? Or is it merely a statement that OLSR is baseline and the IAMANET-team must implement/use it themselves? For example, we have AODV and OLSR working on our system, but surely we don't have any government applications running. We'd like to understand the amount of basic implementation that is needed in order to just get up to the minimal baseline level.

A: The Government intends to supply working code and sample applications.

Q: For the 24-node off-site mini-testbed that will be provided to the third-party evaluator: Will these be returned after the completion of the evaluation? The reason I ask is, as a university, we have a materials-inventory process that might require specifying things differently in the budget if the 24 nodes are not to be returned. I need to know whether I need to work this issue out with my university now.

A: The equipment will be returned.

Q: No specification of the radio platform requirements appears to be given. This has direct bearing on our proposed effort as it could mean the difference between implementing a TDMA overlay ontop of a COTS 802.11, or using a military TDMA radio. I would like to make certain that we are completely allowed to specify this for the purpose of our own research and evaluation.

A: You are allowed to specify the MAC layer for your system.

Q: Is there any specification for the wired network that will be used to test cross-network boundaries? Again, can we choose any wired network that we want?

A: You can define the wired reachback component protocols as needed to be intrinsically assurable.

Q: This is our first experience preparing a DARPA proposal. It would be of immense help if you could outline typical proposal-writing traps and pitfalls; are you permitted to do that?

A: In my experience, here are some prominent proposal-writing traps and pitfalls

- ***Failure to identify a new game-changing, high-payoff, technical idea.*** It is not enough to articulate a worthwhile goal, it is essential to articulate a novel technical insight, perspective, and/or technical approach that has some hope of achieving that goal. Be sure that you have innovative claims, and be sure that you highlight them.
- ***Failure to provide a basis of confidence that your idea is at least plausible.*** It's OK for an approach to be risky, but it is nonetheless important to say why you think it might succeed. Explain preliminary results and analyses if they exist, or your logical argument in the absence of concrete data.
- ***Failure to articulate the hard obstacles that are going to be tackled during Phase 1 and the rest of the program.*** Show that you understand your research agenda by articulating the key challenges defined as quantifiably and/or concretely as possible. Show milestones associated with these challenges in your roadmap. Especially for high-risk projects, it is important to be able to say how and when you will know if you are making progress.
- ***Failure to articulate objective success criteria for your intermediate milestones.*** Ideally you will be able to identify objective and ideally quantitative success criteria for your intermediate milestones that are derived from your technical hypotheses. The greater your level of problem understanding, the more potential there is to generate high-quality milestones.
- ***Failure to explain how and why the program metrics will be achieved as a byproduct of your game-changing technical idea (if it is successful).*** Program metrics rarely reflect the entire gamut of criteria that matter in the real world. A solution that meets the program metrics but fails to address wider concerns is often not a good solution. Meeting the program metrics should be a byproduct of your technical approach, not its only goal.
- ***Failure to address program-specific concerns defined in the solicitation's evaluation criteria.*** DARPA solicitations identify program-specific evaluation criteria that are followed by the source selection committee. Don't forget to think about these criteria when crafting your proposal.
- ***Failure to clarify the relationship between your work and other approaches to the problem.*** It is not enough to claim that your technical idea is new, you should also provide evidence by citing related work. Don't just point out that related work exists; offer some insight as to why your approach is superior. Your choice of related work and your analysis will demonstrate problem understanding and expertise as a byproduct.
- ***Failure to write a proposal specifically for the BAA at hand.*** Proposals about interesting but largely unrelated ideas are submitted with surprising frequency. Odds that they will be selected approach zero; there are other more appropriate ways to bring interesting ideas to DARPA's attention.

Q: I read your feedback on our IAMANET whitepapers and suspect that it was based on some general principles that you applied to every submission. Can you tell us what those principles were?

A: Yes. I did what I could to identify the above traps and pitfalls and generated whitepaper-specific comments about them. In addition, I looked for the following IAMANET-specific issues, which should be considered even if I did not make a specific related comment.

- *Did the whitepaper take full advantage of the “clean slate” opportunity or did it appear to be artificially constrained by limiting assumptions about the existing technology?*
- *Did the whitepaper provide an “analysis of assurability” mapping known threats and vulnerabilities to proposed mechanism for nullifying the threat?*
- *Did the whitepaper describe what it saw as the distinction between the primary and secondary defensive system, and divide its efforts appropriately? Did it explain how the assurable network infrastructure would ultimately simplify and increase the effectiveness of a future secondary defensive system? Note: the objective of Phase 1 is to develop an assurable network infrastructure that eliminates vulnerabilities and prevents attacks. Therefore techniques such as statistical anomaly detection are out of scope per previous Q&A. Also out of scope for Phase 1 are techniques that lose their effectiveness if revealed to the adversary.*
- *Did the whitepaper have something to say about key phase 1 system responses: denying unauthorized traffic by default, authenticating and accounting for all activity on the network, and about tolerating and compartmentalizing byzantine behaviors? If not, was there a specific rationale for the alternative vision?*
- *Did the whitepaper evidence an appropriately broad understanding of the scope of potential byzantine failures and attacks? Note: it is important that proposals recognize that byzantine adversaries can corrupt any node, and effect arbitrary behavior with protocols at every layer.*
- *Did the whitepaper’s approach address integrity, availability, reliability, confidentiality, non-repudiation, and safety?*
- *Did the paper evidence understanding of wireless networking issues?*

Q: We got a lot more negative feedback than positive feedback. Was this really a hint that we should not send in a full proposal?

A: No. If you were encouraged to submit a full proposal, the weaknesses were elaborated in part to help point out matters of form and presentation that would distract reviewers from your substantive ideas.

Q: You did not encourage us from submitting a full proposal. May we submit one anyway?

A: Yes, my encouragement or discouragement is not binding.

Q: You encouraged us to submit a proposal. Does this mean we have a high probability of being selected?

A: No, it means only that I thought your ideas were promising and that with some work you had a reasonable chance of meeting the minimum standards of acceptability in the stated evaluation criteria.

Q: Within our team, we've had a few discussions on the meaning of the exfiltration metric. My current read of the metric is that it is asking us to secure information produced pursuant to running the network and not to secure host data. That is, if the network were using physical locations of nodes for routing, and an attacker were to discover the location of a node by virtue of the routing protocol and then was able to send that data to an accomplice outside the MANET, we would have missed the metric. In contrast, if the attacker were to learn the same information from the SA application and were to send the data to the same accomplice, that would not be covered by the metric.

A: No, your understanding did not match the intent. Both cases matter - what counts is the fact of exfiltration, not where the data came from.

Q: We are confused about the secondary defensive system. Do you have a pre-defined secondary defensive system that you expect our system to accommodate? How are we supposed to develop interfaces to a system we have never seen?

A: There is no predefined secondary defensive system. It will be your responsibility to determine what belongs in the secondary defensive system and what the interface should be.

Questions about scope and specific approaches to the problem

Q: Is anomaly detection of interest as a research focus for this program?

A: Not for Phase 1. For Phase 2, anomaly detection might play a role as an aspect of the secondary defensive system.

Q: There are myriad secure protocols described and proposed in the literature. It might be that the solution for IAMANET already exists, albeit without implementation (proof). So, could the “new” solution be an exercise in fleshing out an “old” idea?

A: Yes, although depending on how “old” the idea is, the evaluation factor concerning novelty may suffer. Be prepared to say why the idea didn’t get traction before, and to argue why it might make sense to apply now.

Q: Have pre-existing security-first-principle technologies been carefully examined for possible leverage? (NSA’s Secure Federated Core X, for example)

A: It is up to the proposer to determine what aspects of previous efforts to reject and what aspects to retain.

Q: Is it desirable to have solutions that allow “authorized list” to depend on the network state as well as on the resources to be used? Will this be tested?

A: It would be up to the proposer to make a case that such a capability is important. Proposers should perform a self-test and evaluation of their solution; the Government plans to perform an independent assessment.

Q: Application software is an inherent component of security. Can we specify what software is part of our system or how software must be written?

A: Proposers may decide to specify an API between applications and the intrinsically assurable infrastructure, and they may require applications to use the API. Keep in mind that all aspects of the proposal, including constraints on application software, will be evaluated for “longevity and impact well beyond the extent of this program challenge”. Therefore, certain kinds of requirements such as an insistence on the use of specific programming languages, or assumptions that applications never contain vulnerabilities in their design and implementation, are not recommended.

Q: Do proposers have to build sample applications? What kinds?

A: Yes, proposers will be responsible for building sample applications. As a convenience, proposers will be provided with relevant demonstration applications. Some envisioned applications include situation awareness, unicast/multicast voice, unicast/multicast video, map viewing, chat, and call-for-fire.

Q: The presence of encryption and other data hiding frequently interferes with secondary defensive systems. Would consideration of network protocol hooks that support trusted intermediaries and other taps be appropriate for phase 1?

A: The design of APIs between the various subsystems is an appropriate consideration for Phase 1.

Q: Would consideration of how to best support particular secondary defensive systems inspection techniques be in scope?

A: Yes, designs that simplify the secondary defensive systems and increase the power of the secondary defensive system are of interest.

Q: Is software diversification of interest as a research focus for this program?

A: No. In some forms software diversification might play a supporting role, but at first glance does not seem to have deep architectural implications, nor does it appear to make a contribution to key desired system responses: accounting for activity on the network, denying unauthorized traffic by default, or tolerating and compartmentalizing byzantine failures.

Topic: fundamental assumptions

Q: What range of channel capacities is targeted by IAMANETs?

A: *DISREGARD ANSWER GIVEN DURING PROPOSER DAY MEETING.* The reconsidered answer is: the widest range of channel capacities that can be made intrinsically assurable. Proposers should specify any assumptions behind their solutions, keeping in mind that certain transition partners foresee continued dependence on low-bandwidth MANETs (HF etc) as well as the availability of high-bandwidth links. Per the evaluation criteria: the “architecture’s scope will be evaluated, insofar as it applies not only to tactical MANETs but also to other networks”. Also considered will be “impediments to transition” and “longevity and impact well beyond the extent of this program challenge”.

Q. Can we assume that the capabilities of a malicious/Byzantine node to be the same as a normal node?

A. Assumptions should always have some basis in the real world, and it is incumbent on the proposer to show problem understanding by detailing the rationale for any of their cornerstone assumptions. For purposes of developing proofs about byzantine robustness, it is often helpful to assume that the adversary’s nodes are even more powerful than normal nodes.

Q. Regarding interoperability with wired networks: should protocol innovations be equally applicable to wireless and wired protocols, or should we bridge to current IP for purposes of backhaul etc?

A. Although the focus of the program is mobile wireless networking, innovations should be applicable to wired networks. Proposers should demonstrate their solutions for a mobile wireless network connected to a wired network. Internet protocol gateways and other forms of IP interoperability are not of interest to this program.

Q: The boundary between primary (direct) trust accountability vs. indirect (anomaly detection) trust may not be “black and white”. What rule of thumb should be used to determine where to draw the boundary?

A: The intent is to keep secret any algorithms or design elements which lose their effectiveness if the adversary knows that they exist and what they are.

Topic: hardware and physical layer technologies

Q. Tamper proof hardware may be needed. Will software level AT be needed as well? How will this be measured? If the Red Team may tamper with the software, will counter-measures be permitted? Is there flexibility in defining tamper proof hardware or selecting hardware purchase in assuming what hardware we need?

A. A research question for this program is whether and how to employ carefully selected tamperproof hardware/software in an assurable MANET design. Since this is merely a research question, any assumptions that are made concerning such a capability must be justified and are open for examination by all parties including the red teams. Designs that rely on such a capability without substantive qualification and justification, and/or that fail to minimize the employment of an anti-tamper capability will be viewed unfavorably. How to build such hardware and software is not the intellectual focus of this program. Selecting such hardware is not the focus of this program. Anti-tamper countermeasures are not of interest to this program.

Q: If we have tamper-proof hardware components with specific authenticity characteristics can the red team substitute them in a way that the defender cannot detect?

A: The issue concerning tamper-proof hardware and how to use it is a research question. The realism, benefits and limitations of a network architecture that relies on such mechanisms must be argued by the proposer. Red teams will work to understand the vulnerabilities of any mechanisms proposed. Spoofing of anti-tamper hardware is a legitimate concern and as such can presumably be modeled by the red team.

Q. Does every node have a HAIPE?

A. No, a “high assurance internet protocol encryptor” is not a requirement in a program that questions whether IP-based architectures are assurable.

Q: What capabilities from the radio system will be available through the API?

A: It is up to the proposer to specify the needed API. Keep in mind that all aspects of the proposal, including constraints on applicable hardware, will be evaluated for “longevity and impact well beyond the extent of this program challenge”.

Q: Are changes to the hardware, such as multiple antenna based node design, to be considered?

A: Hardware changes are not the intellectual focus of this effort; however, architectural designs that exploit both existing and future hardware capabilities are of interest. Keep in mind that all aspects of the proposal, including constraints applicable hardware, will be evaluated for “longevity and impact well beyond the extent of this program challenge”.

Q: What are the examples of baseline hardware that we should consider as part of the Phase I prototyping and evaluation?

A: Hardware is not the intellectual focus of this program. In Phase 1, the hardware will be used for modeling, emulation, and simulation; it can be specified by the proposer.

Q: The physical layer in MANET (whether RF or chips) is vastly neglected in IA as a security mask. Is the physical layer of interest for the IAMANET program?

A: Physical layer security is not a focus of this program, although any architectural design elements needed to exploit the physical layer in support of integrity, availability, confidentiality, etc. are indeed of interest.

Topics: threats and vulnerabilities

Q: Is IAMANET concerned with threats at the application layer or at the lower network layer?

A: The announcement states “The IAMANET program threat model concerns cyberattack in the information domain, including computer worms, pre-inserted malicious code, remote cyber intrusions, exfiltration, protocol exploits, misconfiguration, infrastructure attacks, as well as halting and byzantine failures.” These threats are not layer-specific.

Q: Explain the definition of “lifecycle attack” as used in the metrics slide.

A. A lifecycle attack occurs when the adversary inserts software or hardware vulnerabilities at design, development, production, or deployment time. For purposes of the evaluation, lifecycle attacks can be conservatively modeled by allowing the test team to introduce arbitrary vulnerabilities in selected software or hardware and/or allowing the attacker to have full control of a node during test and evaluation even if that node has not explicitly been infected over the network by a worm or virus. Note that the term “vulnerabilities”, as used here, includes but is not limited to design and programming features allowing systems to be compromised after deployment, and insertion of malicious code into software prior to deployment (e.g., from upstream software suppliers or during the software build process).

Topic: test and evaluation

Q: Will different traffic types and uses have different importance in evaluating throughput performance?

A: Yes. Support for Quality of Service (QoS) is an essential feature of MANETs because bandwidth is limited, and because the need for MANET bandwidth often exceeds availability.

Q: Will the “demonstration MANET” and representative traffic loads involve ground, near-ground and air-based components?

A. Yes.

Q: Currently there is a lack of systematic evaluation methodology for IA evaluation in networks, especially for MANETs. Are evaluation techniques of interest as long as they lead to valid designs?

A: As part of self-test and evaluation, proposers are free to propose any techniques they feel to be informative. The red-team will conduct an independent assessment using tools and techniques of their own choosing.

Q: Will coding flaws in performer software be subject to attack?

A: For purposes of demonstration, yes. However, it is not the intent of the Government to force performers to develop software to unreasonably high standards of quality during Phase 1; the red-teaming analysis will be more heavily focused on what the consequences of such flaws are than whether they exist at all.

Q: Are the internal workings of the red team classified? Will they issue unclassified results or will they issue classified results that we will need to handle with respect to unclassified team members?

A: Not all red-team analyses and findings will be shared with performer teams. If there are classified aspects of those analyses and findings, they will be handled appropriately. In addition, some aspects of the red-team analysis may be covered by ITAR restrictions even if those aspects are not classified.

Topic: other DARPA programs

Q: What are the relationships / synergies/ lessons learned and so on, from CBMANET and ITMANET that apply to IAMANET?

A: All these programs are looking at high-risk/high-payoff ideas. CBMANET shows that novel architectures can offer performance improvements and ITMANET is developing new forms of information theory. Security is not the focus of either program.

Q: Do you anticipate that winning IAMANET proposals will leverage approaches under development in these other programs?

A: Leveraging other DARPA programs is not required.

Topic: proposal preparation and other administrative issues

Q: May a person or organization be a member of multiple proposal teams?

A: Technically yes, but keep in mind that participation on more than one proposal team may dilute your own participation and impact, hinder the synthesis of coherent system ideas, and result in a lower-quality proposal to the detriment of all concerned.

Q: How much do you want to pay for assurance? There is a difference between “assurable” and “assured”; for instance, the code for an “assured” network may have been evaluated line-by-line by teams of mathematicians.

A: As little as possible, but no less. Proposals will be evaluated on cost and schedule reasonableness and realism. Keep in mind that the purpose of this DARPA program is to prove the value of some new ideas, not to harden a product. Ideally your techniques

minimize the need for laborious human effort both in the research phase and during subsequent transition.

Q: Does an RA proposal have to have industry participation?

A: No. In fact excessive or gratuitous industry participation in an RA proposal could lead to a determination that the proposal is non-responsive. Please read the proposer information pamphlet carefully.

Q: A proposer day slide implied that red team attacks in Phase I may be classified. Please explain.

A: True. Red teams are not obligated to reveal their attack techniques to performers, and may choose to consider techniques not known to the general public.

Q: Can a proposer respond to Phase 1 only?

A: All proposals are for Phase 1 only, however, proposals that do not outline an intent and a means to participate in Phase 2 will be considered non-responsive.

Q: Is there a limit on the number faculty (early career or otherwise) that may be involved in IAMANET?

A: No. The duration and budget of Phase 1 are up to the performer and should be necessary and sufficient to meet the research objectives. Keep in mind that one of the evaluation criteria is “cost and schedule reasonableness and realism”.

Q: What travel is expected for the Phase I prototype?

A: PI meetings as detailed in the proposer information pamphlet, plus any travel needed to support the 24-node testbed at the red-team facility.

SN 07-30, Intrinsically Assurable Mobile Ad-Hoc Networking (IAMANET)

Questions and Answers

Updated 4/12/2007

Q: Is the IAMANET RA an opportunity for individual young faculty to submit individual proposals to the program, or would it involve forming larger teams?

A: DARPA plans to request proposals for the full scope of Phase 1 research and prototyping (i.e., an end-to-end system designed by a team of multidisciplinary research organizations, plus an integrator for coordination and implementation support). Proposals addressing only individual component-level technologies will be considered non-compliant with the requirements of this solicitation.

Q: How does industry involvement factor into the early career research announcement (RA)?

A: See the "Eligible proposers" section of the Proposer Day Announcement. "It is recognized that the proposal team may utilize non-proposal team members as sub awardees for certain aspects of the proposed work. However, a majority of the research effort and leadership of the technical direction must reside with the Early-Career Investigators." Proposals with excessive or gratuitous industry involvement will be considered noncompliant.

Q: How should I parse the following sentence: "However, a majority of the research effort and leadership of the technical direction must reside with the Early-Career Investigators"?

A: Apologies for the ambiguity. The correct parsing is "However, a (majority of the research effort) and (leadership of the technical direction) must reside with the Early-Career Investigators." To be clearer: all responsibility for the project's technical direction must reside with the Early-Career investigators.

Q: Would it be considered a conflict of interest for a company to go in on an RA as well as submit its own BAA response?

A: Not according to legal definitions of "conflict of interest" that I am aware of. However see the restrictions on industry involvement in the RA as specified in the "Eligible proposers" section.

Q: To what level is citizenship needed?

A: Citizenship becomes an issue with respect to U.S. export control laws. Phase 1 is not expected to involve any ITAR material. Notwithstanding, the following provision will be incorporated into any resultant contract or grant:

(1) The contractor shall comply with all U. S. export control laws and regulations, including the International Traffic in Arms Regulations (ITAR), 22 DFR Parts 120 through 130, and the Export Administration Regulations (EAR), 15 CFR Parts 730 through 799, in the performance of this contract. In the absence of available license exemptions/exceptions, the Contractor shall be responsible for obtaining the appropriate licenses or other approvals, for obtaining the appropriate licenses or other approvals, if required, for exports

of hardware, technical data, and software, or for the provision of technical assistance.

(2) The Contractor shall be responsible for obtaining export licenses, if required, before utilizing foreign persons in the performance of this contract, including instances where the work is to be performed on-site at any Government installation, including installations within the United States, where the foreign person will have access to export-controlled technical data or software.

(3) The Contractor shall be responsible for all regulatory record keeping requirements associated with the use of licenses and license exemptions/exceptions.

(4) The Contractor shall be responsible for ensuring that the provisions of this clause apply to its subcontractors.

Q: How does citizenship factor into the decision process?

A: It does not. However, "proposers to Phase 1 must intend to participate in subsequent phases, and will be expected to outline the means whereby they will be able to continue work in subsequent phases." See "eligible proposers" and "security considerations for eligible proposers" for more details. Proposals not meeting this standard will be considered non-compliant.